arXiv:1402.5524v1 [cs.CR] 22 Feb 2014

The Melbourne Shuffle: Improving Oblivious Storage in the Cloud

Olga Ohrimenko[1] Michael T. Goodrich[2] Roberto Tamassia[3] Eli Upfal[3]

[1]Microsoft Research, [2]University of California, Irvine, [3]Brown University

Abstract

We present a simple, efficient, and secure data-oblivious randomized shuffle algorithm. This is the first secure data-oblivious shuffle that is not based on sorting. Our method can be used to improve previous oblivious storage solutions for network-based outsourcing of data.

1 Introduction

One of the unmistakable recent trends in networked computation and distributed information management is that of *cloud storage* (e.g., see [15]), whereby users outsource data to external servers that manage and provide access to their data. Such services relieve users from the burden of backing up and having to maintain access to their data across multiple computing platforms, but, in return, such services also introduce privacy concerns. For instance, it is likely that cloud storage providers will want to perform data mining on user data, and it is also possible that such data will be subject to government searches. Thus, there is a need for algorithmic solutions that preserve the desirable properties of cloud storage while also providing privacy protection for user data.

Of course, users can encrypt data they outsource to the cloud, but this alone is not sufficient to achieve privacy protection, because the data access patterns that users exhibit can reveal information about the content of their data (e.g., see [4, 14]). Therefore, there has been considerable amount of recent research on algorithms for *data-oblivious algorithms and storage*, which hide data access patterns for cloud-based network data management solutions (e.g., see [9, 10, 11, 12, 13, 18, 21, 22, 23, 25, 26]). Such solutions typically work by obfuscating a sequence of data accesses intended by a client by simulating it with the one that appears indistinguishable from a random sequence of data accesses. Often, such a simulation involves mixing the intended (real) accesses with a sequence of random "dummy" accesses. In addition, so as to never access the same address twice (which would reveal a correlation), such obscuring simulations also involve continually moving items around in the server's memory space. For this reason, the "inner-loop" computation required by such simulations is a *data-oblivious shuffling* operation, which moves a set of items to random locations in fashion that disallows the server to correlate the previous locations of items with their new locations. This inner-loop process requires putting items in new locations that are independent of their old locations while hiding the correlations between the two.

The most common way this inner-loop shuffling is implemented is, however, computationally expensive, since it involves assigning random (or pseudo-random) indices to items and then performing a data-oblivious sorting of these index-item pairs. Examples of such oblivious sorting algorithms include Batcher's sorting network [3], which requires $O(n(\log n)^2)$ I/Os to sort data of size n, or the AKS [1] or Zig-zag sorting [8] networks, which use $O(n \log n)$ I/Os, but with large constant factors that restrict their practicality. These algorithms are used in oblivious storage solutions by having a client use the server as an external memory, with the I/Os directing the client to issue commands to move items from the server to the client's private memory and from the client's private memory to the server. Though these solutions achieve a desired privacy

level, they are expensive in their (amortized) access overhead time and also in their monetary cost when one considers a client outsourcing large volumes of data and accessing it from a cloud server that charges per every data request.

In this paper, therefore, we are interested in algorithmic improvements for oblivious storage solutions, in terms of their conceptual complexity, constant factors, and monetary costs. For instance, since cloud-storage servers typically charge users for each memory access but have fairly large bounds on the size of the messages for such I/Os, we allow for messages to have modest sizes, such as $O(\sqrt{n})$ for a storage of size n. This necessarily also implies that the client has an equally modest-sized private memory, in which to send and receive such messages (and also in which to perform internal swaps of data items away from the prying eyes of the server). Our goal in this research is to take advantage of such frameworks to replace data-oblivious sorting with simple oblivious data shuffling for the sake of providing simple, efficient, and cheap outsourced data management. Our framework, therefore, involves designing (or modifying) oblivious storage simulation algorithms where a client stores n items at the server and is allowed to issue a sequence of I/Os, each of which is a batch of reads and writes for the server's memory, for reasonable assumptions on message size and private memory size.

Related Work. A *shuffle* is an algorithm for rearranging an array to achieve a random permutation of its elements. Early shuffle methods were motivated by the problem of shuffling a deck of cards.

Classic card shuffle methods (e.g., Knuth (or Fisher-Yates) [17], the riffle shuffle [2], Thorp shuffle [24]) are not data-oblivious, however, as anyone observing card swaps or riffles (interleaving two subdecks) of such methods can learn the final output permutation. In ICALP 2012, Goodrich and Mitzenmacher [10] showed that one can, in fact, shuffle a deck of n cards and guarantee that an observer cannot find a particular card in the output permutation with probability better than $O(1/n)$. However, this algorithm is not an effective shuffle for our purposes, since the output permutations produced by the algorithm are not all equally likely and there may be dependencies between large groups of cards that could be leaked. Most other existing efficient data-oblivious shuffling methods assign random values to the elements of the array and use a data-oblivious algorithm to sort the array according to these values.

Our Oblivious Shuffling Results. Our Melbourne shuffle[1] algorithm is instead the first data-oblivious shuffle method that is not based on a data-oblivious sorting algorithm. In Table 1, we compare the Melbourne shuffle, showing that it outperforms sorting-based shuffle methods.

Table 1: Comparison of data-oblivious sorting and shuffle algorithms over n items.

	Randomized	Private Memory	Message Size	External Memory	I/Os
Batcher's network [3]		$O(1)$	$O(1)$	$O(n)$	$O(n(\log n)^2)$
Batcher's network I		$O(\sqrt{n})$	$O(\sqrt{n})$	$O(n)$	$O(\sqrt{n}(\log n)^2)$
Batcher's network II [9]		$O(\sqrt{n})$	$O(\sqrt[s]{n})$	$O(n)$	$O(n^{7/8})$
AKS [1], Zig-zag sort [8]		$O(1)$	$O(1)$	$O(n)$	$O(n \log n)$
Randomized shellsort [7]	✓	$O(\sqrt{n})$	$O(\sqrt{n})$	$O(n)$	$O(\sqrt{n} \log n)$
Melbourne shuffle	✓	$O(\sqrt{n})$	$O(\sqrt{n})$	$O(n)$	$O(\sqrt{n})$
Melbourne shuffle ($c \geq 3$)	✓	$O(\sqrt[c]{n})$	$O(\sqrt[c]{n})$	$O(n)$	$O(c\sqrt[c]{n^{c-1}})$

Improved Oblivious Storage. Oblivious storage and oblivious RAM (ORAM) simulation solutions aim at minimizing the *access overhead*, which is the amortized number of I/Os executed to perform a single stor-

[1]The name of our algorithm is inspired by a ``shuffle'' dance technique.

age access request while keeping reasonable assumptions about the size of private memory and of messages exchanged between the client and the server, e.g., sublinear in the size of outsourced memory n. In seminal work motivated by a software protection application, Goldreich and Ostrovsky [5, 6] give two oblivious storage solutions for a client with $O(1)$ private memory size: the square root method with $O(\sqrt{n})$ overhead and the hierarchical method with $O((\log n)^3)$ overhead. The hierarchical method was recently extended using techniques such as Bloom filters [25, 26] and cuckoo hash tables [9, 13, 18, 20]. E.g., the $\log n$-hierarchical solution of [13] uses $O(\sqrt[d]{n})$ temporary memory and achieves $O(\log n)$ access overhead, for $d \geq 2$. In [18] a similar method achieves $O((\log n)^2 / \log\log n)$ overhead with $O(1)$ private memory.

All the above oblivious storage solutions rely on a periodic data-oblivious shuffle of the server storage, a task done using data-oblivious sorting. This is the most expensive step of such solutions, but since it happens only after a certain number of requests, it can be amortized. Otherwise, one can use techniques of [11] to deamortize these solutions to bring the worst case access overhead to be the same as the average case. Thus, we can use our Melbourne shuffle to implement the shuffle steps of these algorithms.

Other oblivious storage solutions in [21, 22, 23] allow the user to have $o(n)$ private memory and then by applying the same solution recursively on this private memory bring it to $o(1)$ and adding a $\log n$ overhead. For example, Path ORAM [23] uses $O(\log n)$ (stateful) private memory and has $O((\log n)^2)$ access overhead.

As shown in Table 2, oblivious storage solutions based on our Melbourne shuffle are efficient and practical.

Table 2: Comparison of oblivious storage solutions for n items.

	Private Memory	Message Size	External Memory	Access Overhead
SquareRoot [6]	$O(1)$	$O(1)$	$O(n)$	$O(\sqrt{n})$
Path ORAM [23]	$O(\log n)$	$O(\log n)$	$O(n)$	$O((\log n)^2)$
Bucket Hash Hierarchical [6]	$O(1)$	$O(1)$	$O(n \log n)$	$O((\log n)^3)$
Cuckoo Hash Hierarchical [13] ($d \geq 2$)	$O(\sqrt[d]{n})$	$O(\sqrt[d]{n})$	$O(n)$	$O(\log n)$
SquareRoot with Melbourne shuffle	$O(\sqrt{n})$	$O(\sqrt{n})$	$O(n)$	$O(1)$
Hierarchical with Melbourne shuffle ($c \geq 3$)	$O(\sqrt[c]{n})$	$O(\sqrt[c]{n})$	$O(n)$	$O(c \log n)$
	$O(\sqrt[c]{n} \log n)$	$O(\sqrt[c]{n} \log n)$	$O(n)$	$O(c)$

2 Preliminaries

2.1 Cryptographic Primitives

We analyze the security of cryptographic primitives and our protocol in terms of the probability of success for an adversary in breaking them. Let k be a security parameter. We consider a probabilistic adversary $\mathcal{A}$ whose running time is polynomial in k. We say that a scheme is secure if for every probabilistic polynomial time (in k) (PPT) adversary $\mathcal{A}$, the probability of breaking the scheme is at most some *negligible function* $\mathsf{negl}(k)$, i.e., a function such that $\mathsf{negl}(k) < 1/|\mathsf{poly}(k)|$ for every polynomial $\mathsf{poly}(k)$.

CPA Secure Encryption We use a *symmetric encryption scheme* $(\mathsf{Enc}_{\mathsf{key}}, \mathsf{Dec}_{\mathsf{key}})$ where $\mathsf{key} \leftarrow \{0,1\}^k$. We require this scheme to be secure against the chosen-ciphertext attack (CPA) for multiple messages [16]. During this attack an adversary $\mathcal{A}$ is allowed to make queries to oracles $\mathsf{Enc}_{\mathsf{key}}$ and $\mathsf{Dec}_{\mathsf{key}}$ on a polynomial number of sequences of l messages of his choice. After this "warm-up" phase, $\mathcal{A}$ comes up with two sequences, M_0 and M_1, of l messages and gives them to a challenger. The challenger secretly picks a bit b and calls $\mathsf{Enc}_{\mathsf{key}}$ on each message of sequence M_b. Let C be the sequence of ciphertexts that correspond

to M_b. The challenger gives C to $\mathcal{A}$ who continues querying $\mathsf{Enc}_{\mathsf{key}}$ and $\mathsf{Dec}_{\mathsf{key}}$ on any sequence of ciphertexts except those in C for a polynomial number of times. Finally, the adversary's task is to guess bit b. We call the above game *Enc-IND-CPA* and say that $\mathcal{A}$ wins the game if he correctly guesses b. Then, $(\mathsf{Enc}_{\mathsf{key}}, \mathsf{Dec}_{\mathsf{key}})$ is said to be secure if for all PPT adversaries, the probability of winning game Enc-IND-CPA is at most $1/2 + \mathsf{negl}(k)$. We omit using key when referring to $(\mathsf{Enc}, \mathsf{Dec})$. For an intuition behind Enc-IND-CPA secure encryption scheme, consider encrypting a message padded with a different random nonce each time it is encrypted. Hence, re-encryptions of the same plaintext look different with very high probability.

Pseudo-Random Permutation (PRP) Consider an array of n elements that we wish to randomly rearrange and let $D = [1, n]$ be the set of indices of A. We use a family of efficiently computable *pseudo-random permutations* (PRPs) $\Pi_{\mathsf{seed}} : D \to D$, keyed using a seed from the set $\mathsf{Seeds}(\Pi) = \{0,1\}^k$ [16]. In order to pick a permutation, one picks a random seed from $\mathsf{Seeds}(\Pi)$ and stores it privately. Hence, whenever we refer to choosing a permutation, we refer to picking a new $\mathsf{seed} \in \mathsf{Seeds}(\Pi)$. Once the seed is fixed, we can evaluate Π_{seed} on a given index $x \in D$ via $\Pi_{\mathsf{seed}}(x)$.

The security of a family of PRPs is defined by comparing the behavior of a PPT adversary when he is given a truly random permutation versus a pseudo-random permutation picked using a random seed. Formally, let $\mathcal{R}$ be the set of all permutations over the domain D. A family of PRPs Π is secure if for every probabilistic polynomial time adversary $\mathcal{A}$, the probability of distinguishing between $\mathsf{r} \xleftarrow{\$} \mathcal{R}$ and Π_{seed}, where $\mathsf{seed} \xleftarrow{\$} \mathsf{Seeds}(\Pi)$, is $1/2 + \mathsf{negl}(k)$.

Given an array A of n (key,value) pairs (x, v) where $x \in [1, n]$, we denote the permutation π of A as $B = \pi(A)$, where $\pi = \Pi_{\mathsf{seed}}$ and $B[x] = A[\pi(x)], \forall x \in [1, n]$. We will use the same notation when A and B are encrypted. We refer to the original permutation of A, as permutation π_0. For A, sorted using x, π_0 is the identity.

2.2 Storage Model

We consider a cloud storage model where a client stores a dataset at a server while keeping a small amount of data in private memory. For simplicity, we assume that the dataset is an array of elements of equal size. The client encrypts each element and stores the elements at the server according to a pseudo-random permutation. The encryption key and the seed of the permutation are kept private by the client and are not revealed to the server.

Client Private Memory We assume that the client has access to a small private memory, M, which is comprised of permanent storage and scratch space. The permanent storage includes the encryption key and the current seed of the permutation the client is using, which together is of size $O(1)$. The rest of M is used as a scratch space while performing operations on the remote storage and is not needed in between operations. We require the size of the scratch space to be sublinear in n. Depending on the algorithm, we will use a private storage of size $\sqrt[c]{n} \log n$ or $\sqrt[c]{n}$, for an arbitrary integer $c \geq 2$.

Since the client can store a small number of elements at a time, we assume that he does not try to request from the server more than he can fit and process in M. Let the *message size*, denoted $\mathsf{msgSize}$, be the maximum elements that can be exchanged by the client and server in one operation. We have that $\mathsf{msgSize}$ should be less than the size of the scratch space.

Server Memory The server supports the following operations on an array S.

- $\mathsf{get}(S, \mathsf{loc})$: return element stored at a location loc in S.

- $\mathsf{put}(S, \mathsf{loc}, e)$: put element e to a location loc in S.
- $\mathsf{getRange}(S, \mathsf{loc}, \ell)$: return an array a with elements at locations $\mathsf{loc}, \ldots, \mathsf{loc} + \ell - 1$ in S, where $\ell \leq \mathsf{msgSize}$.
- $\mathsf{putRange}(S, \mathsf{loc}, a)$: write elements in array a to locations $\mathsf{loc}, \ldots, \mathsf{loc} + |a| - 1$ in S, where $|a| \leq \mathsf{msgSize}$.
- $\mathsf{getRangeDist}(S, \langle \mathsf{loc}_1, \ldots, \mathsf{loc}_c, \rangle, \langle \ell_1, \ldots, \ell_c, \rangle)$: return an array a with elements at locations $\mathsf{loc}_i, \ldots, \mathsf{loc}_i + \ell_i - 1$ in $S, \forall i \in [1, c]$, where $\sum \ell_i \leq \mathsf{msgSize}$.
- $\mathsf{putRangeDist}(S, \langle \mathsf{loc}_1, \ldots, \mathsf{loc}_c \rangle, \langle a_1, \ldots, a_c \rangle)$: write elements in each array a_i to locations $\mathsf{loc}_i, \ldots, \mathsf{loc}_i + |a_i| - 1$ in S, where $\sum |a_i| \leq \mathsf{msgSize}$.

Note that the number of elements in getRange and putRange is limited by the maximum number of elements that can be exchanged between the client and the server in one operation.

We assume that the server can perform operations get and put in constant time and operations getRange, putRange, getRangeDist and putRangeDist in time proportional to the number of elements read or written, but each operation takes one I/O.

Definition 2.1 (Metadata): *The name of the array S, its size, location i, l, and the size of a are referred to as the* metadata *of a* getRange *or a* putRange *call. Similarly for* getRangeDist *and* putRangeDist*, locations $\langle \mathsf{loc}_1, \ldots, \mathsf{loc}_c \rangle$, $\langle \ell_1, \ldots, \ell_c, \rangle$ and sizes of $a_1, \ldots, a_c$ are referred as metadata as well.*

3 Oblivious Shuffle Model

In this section, we introduce a formal model for the oblivious shuffle of an array.

3.1 Model

Definition 3.1 (Shuffle): *We define a* shuffle $\mathcal{S}$ *as a pair of algorithms* $(\mathsf{Setup}, \mathsf{Shuffle})$*, as follows.*

- $(s, S) \leftarrow \mathsf{Setup}(1^k)$ *Given security parameter k, run the key generation algorithm for a symmetric encryption scheme* $(\mathsf{Enc}, \mathsf{Dec})$ *and store the key in secret state s. Also, allocate an auxiliary datastore S.*
- $(\mathsf{Enc}(\pi(A)), \alpha) \leftarrow \mathsf{Shuffle}(s, S, A, \pi)$ *Given secret state s, auxiliary data store S, an array input A, and a permutation π, return (1) the encryption of the permutation of A according to π; (2) a transcript α of the operations that transform* $\mathsf{Enc}(A)$ *to* $\mathsf{Enc}(\pi(\mathsf{A}))$ *using auxiliary space S.*

Transcript α is a sequence of l (request, response) pairs $\langle (r_1, g_1), \ldots, (r_l, g_l) \rangle$ that capture the evolution of the datastore via intermediate states $S_1, S_2, \ldots, S_{l+1}$. An invariant on each intermediate state is to store an encryption of some permutation of A along with any auxiliary data. For example S_1 contains $\mathsf{Enc}(A)$ *and S_l contains* $\mathsf{Enc}(\pi(A))$*. Setting $S_1 \leftarrow \{\mathsf{Enc}(A), S\}$, $s_0 \leftarrow s$, $g_0 \leftarrow \perp$ define the relationship between r_i and g_i as:*

$$\langle\ (s_i, r_i) \leftarrow \mathsf{GenRequest}(s_{i-1}, g_{i-1}),\ \ (S_{i+1}, g_i) \leftarrow \mathsf{GenResponse}(S_i, r_i)\ \rangle.$$

Operations GenRequest *and* GenResponse *generate a request r_i and a corresponding response g_i and are defined as follows:*

- $(s_i, r_i) \leftarrow \mathsf{GenRequest}(s_{i-1}, g_{i-1})$ *Perform a computation based on a substructure of* S_{i-1}, g_{i-1}, *and generate next request to* S_i, r_i.
- $(S_{i+1}, g_i) \leftarrow \mathsf{GenResponse}(S_i, r_i)$ *Generate the response to request* r_i *on* S_i*:* S_{i+1} *is the datastore* S_i *updated according to* r_i *and* g_i *is the response to* r_i *with respect to* S_i*. For example, if* r_i *is a get request, then* $S_{i+1} = S_i$ *and* g_i *is the requested item. Also, if* r_i *is a put request, then* g_i *is empty.*

The private state s *is updated if needed after every request.*

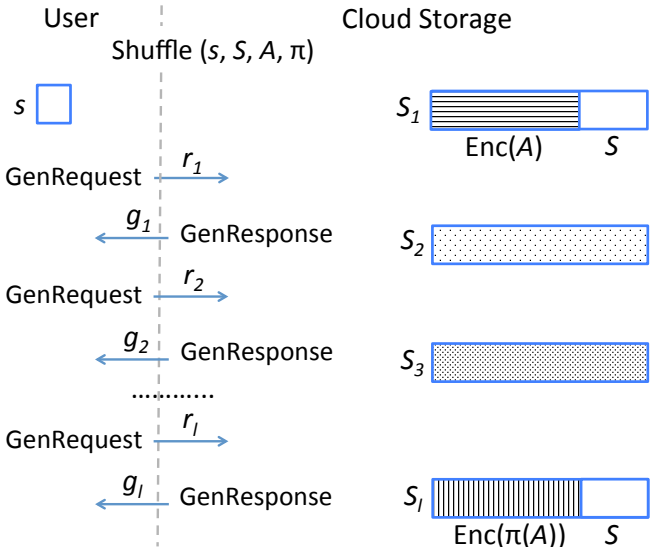

Figure 1: Illustration of a shuffle $\mathcal{S}$ executed by the client and the cloud storage server (Section 3.1).

In our cloud storage model, a shuffle $\mathcal{S}$ is a distributed computation executed by the user and the server. The user runs the Setup algorithm to generate the encryption key and requests the server to allocate some space. He then runs the Shuffle algorithm by accessing S through the server, that is, issuing requests to the server using GenRequest. The set of possible requests is defined by the storage model supported by the server. In our case this set is {get, put, getRange, putRange, getRangeDist, putRangeDist} (see Section 2.2). For every request r_i, the server executes GenResponse, locally updating S for put requests and returning to the user the queried items for get requests. (See Figure 1 for an illustration.)

3.2 Security

We capture the security of a shuffle $\mathcal{S}$ against a curious server in the cloud storage model as a game, *Shuffle-IND*, between $\mathcal{S}$ and a probabilistic polynomial-time bounded (PPT) adversary $\mathcal{A}$. In this game, the inputs and outputs of $\mathcal{S}$ that are revealed to the server in the cloud storage model are also revealed to $\mathcal{A}$. However, the secret state s kept by the client, any updates to it and computations inside of GenRequest are kept private, since in the cloud model they are also hidden and happen on the user side.

The game starts with S running Setup once, allocating at the server space to be used in subsequent computations. $\mathcal{A}$ then tries to "learn" how $\mathcal{S}$ performs the shuffle on a sequence of m_1 input arrays and permutations picked by $\mathcal{A}$. Based on what $\mathcal{A}$ learns, she picks two challenges (A_0, τ_0) and (A_1, τ_1) each consisting of a data array to be permuted using a corresponding permutation. $\mathcal{S}$ secretly picks one pair and performs the shuffle according to it. The adversary is then allowed to observe $\mathcal{S}$ shuffling another sequence of m_2 (input, permutation) pairs, also picked by $\mathcal{A}$. Finally, $\mathcal{A}$ has to guess which challenge pair

(input, permutation) $\mathcal{S}$ picked to shuffle. Note that at any time, $\mathcal{A}$ can ask $\mathcal{S}$ to perform a shuffle on any combination of A_0 or A_1 and permutations τ_0 or τ_1.

We now give a formal definition of the game.

Definition 3.2 (Shuffle-IND): *Let A be an input array of size n picked by a PPT adversary $\mathcal{A}$. $\mathcal{A}$ and $\mathcal{S}$ engage in the following game.*

$\mathcal{S}$**:** $(s, S) \leftarrow \mathsf{Setup}(1^k)$.

for $j \in \{1, \ldots, m_1\}$, where m_1 is $\mathsf{poly}(k)$:

$\mathcal{A}$: Pick array B_j and a permutation ρ_j.

$\mathcal{S}$: Execute $(O_j, \alpha) \leftarrow \mathsf{Shuffle}(s, S, B_j, \rho_j)$. Reveal O_j and α to $\mathcal{A}$.

$\mathcal{A}$: Pick (A_0, τ_0) and (A_1, τ_1) of the same length.

$\mathcal{S}$: Pick a secret bit b and execute $(O, \alpha) \leftarrow \mathsf{Shuffle}(s, S, A_b, \tau_b)$. Reveal O and α to $\mathcal{A}$.

for $j \in \{m_1 + 1, \ldots, m_1 + m_2\}$, where m_2 is $\mathsf{poly}(k)$:

$\mathcal{A}$: Pick an encrypted array B_j and a permutation ρ_j.

$\mathcal{S}$: Execute $(O_j, \alpha) \leftarrow \mathsf{Shuffle}(s, S, B_j, \rho_j)$. Reveal O_j and α to $\mathcal{A}$.

$\mathcal{A}$: Output bit b'.

The adversary wins the game if $b = b'$.

Using the Shuffle-IND game, we now define an *oblivious shuffle*.

Definition 3.3 (Oblivious Shuffle): *Let k be the security parameter and n be a polynomial in k. $\mathcal{S}$ is an oblivious shuffle over n items if for every probabilistic adversary $\mathcal{A}$ running in time polynomial in k, the probability of winning the Shuffle-IND game, $\Pr[b = b']$, satisfies*

$$\Pr[b = b'] \leq \frac{1}{2} + \mathsf{negl}(k).$$

3.3 Performance

We measure the performance of an oblivious shuffle of an array of n items A using the following parameters:

Number of Requests: the number of calls the user and the server make to each other while performing the shuffle. In the protocol, the number of requests is expressed using l, the length of the transcript α. This parameter measures the efficiency of a shuffle algorithm.

Message Size: the maximum number of items that can be sent between the user and the server in a single operation, i.e., the number of items sent in a request r_i or a response g_i.

User Private Memory: the size of user's private memory s measured in terms of the number of items of A that can be stored temporarily plus the space required to store an encryption key and a seed for a pseudo-random permutation. The key and the seed are stored in the stateful part of s which is of a constant size, while the space required to store the items is the scratch space required only during the execution of the shuffle algorithm and is erased afterwards.

User Computation: the amount of computation the user performs during the Shuffle algorithm, i.e., computation inside of $\mathsf{GenRequest}(s_i, g_{i-1})$.

Server Storage: additional space required at the server besides storing n encrypted items of A. This is captured by the datastore S in the protocol.

Server Computation: the computation performed by the server during $\mathsf{GenResponse}(S_i, r_i)$.

In the cloud storage model, we wish to devise efficient shuffle solutions that consider realistic assumptions about the user. For example, a user private memory of size n leads to a trivial and efficient solution where the user can download the data from the server, shuffle and re-encrypt the items, and send them back. Instead we wish to construct solutions that add a small overhead over a non-oblivious solution in terms of the number of requests and client computation, but assume private memory and message size sublinear in n. We will also assume that the server and the user can exchange more than one item in one message or call. The size of each message is limited by the size of user's memory since this is how many items she can process at a time.

Finally, we consider current cloud storage providers that store user data and can efficiently retrieve and write small amounts of data as requested by the user. Such storage providers charge their users according to a pay-per-use model. Hence, we wish to limit space requirements of our solutions.

4 The Melbourne Shuffle

In this section, we present a basic version of our Melbourne shuffle algorithm. An optimized version is given in the next section. The basic Melbourne shuffle uses private memory and messages of size $O(\sqrt{n}\log n)$, $O(n \log n)$ server storage and processes in $O(\sqrt{n})$ requests. The optimized Melbourne shuffle has a smaller message and memory overhead, and requires constant number more accesses to the server. In particular, it relies on private memory and messages of size $O(\sqrt{n})$, $O(n)$ server storage and $O(\sqrt{n})$ accesses.

An important ingredient of our solution is probabilistic encryption. Everything stored at the server is encrypted and every time an item is read from the server, the user decrypts it, re-encrypts it and writes it back. Since we use CPA-secure encryption, the ciphertexts produced for the same item always look different and, hence, the server, aka the adversary, cannot tell whether the ciphertexts correspond to the same item or not.

The goal of our oblivious shuffle is to reveal to the adversary only information that she would expect to see in a random permutation with very high probability. For example, even for a secret permutation picked uniformly at random, the adversary can guess with probability $1/n$ that the first element of the input array of size n appears in some location i of the output permutation. Continuing with this intuition, suppose we split the input array of size n into $\sqrt{n}$ buckets where every bucket has $\sqrt{n}$ items, and similarly for the output permutation. In this case, using the analysis of the balls-and-bins model, the adversary can guess that with high probability, each bucket in the output permutation has $O(\log n)$ elements from any particular bucket of the input array.

We build on the observation above and move elements from input buckets to output buckets by imitating the balls-and-bins process. That is, if the size of the input bucket is the same as the number of output buckets, we place $O(\log n)$ elements of every input bucket in every output bucket. If the number of elements in a bucket is much larger than the number of output buckets, i.e., the number of output buckets is $n^{1/c'}$ while input bucket has $n^{1/c}$ items, for constants c and c' s.t. $c' > c$, then we move $O(n^{1/c-1/c'})$ items to every output bucket.

The reader may have noticed that in the first example above, elements of an input bucket of size $\sqrt{n}$ are placed in $\sqrt{n}$ output buckets in batches of $O(\log n)$ items. What are the additional items? These additional

items are referred to as *dummy items*. A dummy item is a real item with a fake key and some nonce value such that the size of the dummy and real item are equal. Moreover, since all the data is re-encrypted every time it is written to the server, the server cannot tell which items in a batch are real and which are dummy.

4.1 Overview

We assume that each element in the input array, A, is a key-value pair (x, v) for every $x \in D$. The algorithm proceeds in two phases: *distribution* and *clean-up*. For each phase, the data store S is split in several logical subparts: I, T and O. I is an array containing n encrypted items of the input A permuted according to some permutation π_0 (initially, π_0 is the identity). T is an encrypted temporary array used during the shuffle; finally, after the shuffle is done O contains the output of the shuffle, i.e., re-encrypted items of I permuted according to π. If the shuffle needs to be executed again, the user sets $I \leftarrow O$ and $\pi_0 \leftarrow \pi$. We further divide each subpart of S in buckets of equal size. The number of buckets and how it effects the runtime of the algorithm will be determined later.

During the distribution phase items of every bucket of I along with some dummy items are re-encrypted and distributed equally among buckets of T. Here, the distribution of item (x, v) is done according to its final location $\pi(x)$ in O. After the distribution phase the intermediate array T contains real and dummy items. Moreover, the items appear in correct buckets but not in correct positions within each bucket. The clean-up phase remedies this by reading one bucket at a time, removing dummy items, distributing the real items correctly within the bucket and writing the bucket to O.

The distribution phase alone cannot produce every possible permutation since the number of items sent from a bucket of I to a bucket of T is limited. E.g., the identity permutation cannot be achieved. To rectify this, we execute two shuffle passes. First, for a permutation π_1 picked uniformly at random and then for the desired permutation π. Although this framework still allows failures, our algorithm can produce every permutation, failing with very small probability independent of the desired permutation π.

4.2 Algorithm

The complete shuffle algorithm $\mathsf{shuffle}(I, \pi, O)$ is shown in Algorithm 1 where I is the encryption of the input array A, π is the desired permutation and the last argument is the output array where the algorithm is expected to put an encryption of $\pi(A)$. We omit the Setup from the discussion since it is trivial: the client simply runs a key generation to setup a secure encryption scheme and a seed generator for pseudo random permutations.

The algorithm makes two calls to shuffle_pass (Algorithm 2), first for a random permutation π_1 and then for the desired permutation π. We proceed with the description of $\mathsf{shuffle_pass}(I, T, \rho, O)$ where I and O are defined as in shuffle, T is a temporary array and ρ is the desired permutation. We use the convention of giving arrays I, T and O as inputs to the shuffle pass algorithm for the ease of explanation. In the cloud storage scenario that we consider here, one simply specifies the location where these arrays are stored remotely, e.g., the name of a file and a location within it. Given an input array of size n, this method has messages and client's private memory of size $O(\sqrt{n} \log n)$ and server memory of size $O(n \log n)$. These user and server memory requirements are temporary and are reduced to $O(1)$ and n, respectively, when the shuffle is finished. As mentioned before, method shuffle_pass is split into a distribution phase and a clean-up phase.

Distribution Phase The distribution phase of method shuffle_pass (Algorithm 2), shown in Figure 2, imitates throwing balls into bins by putting elements from every bucket of I to every bucket of T according to the permutation ρ. In particular, a batch of $p \log n$ encrypted elements from every bucket of I is put

Algorithm 1 The complete Melbourne shuffle algorithm, shuffle(I, π, O), where the user can read and store in private memory M up to $\sqrt{n} \times p \log n$ elements, $p \geq$ e.

$\boldsymbol{I}$: array of n encrypted elements (x, v); $\boldsymbol{\pi}$: permutation; $\boldsymbol{O}$: permutation of I according to π, where every element is re-encrypted.

1: Let π_1 be a random permutation
2: Let T be an empty array of size $n \times p \log n$ stored remotely
3: shuffle_pass(I, T, π_1, O)
4: $I \leftarrow O$
5: shuffle_pass(I, T, π, O)

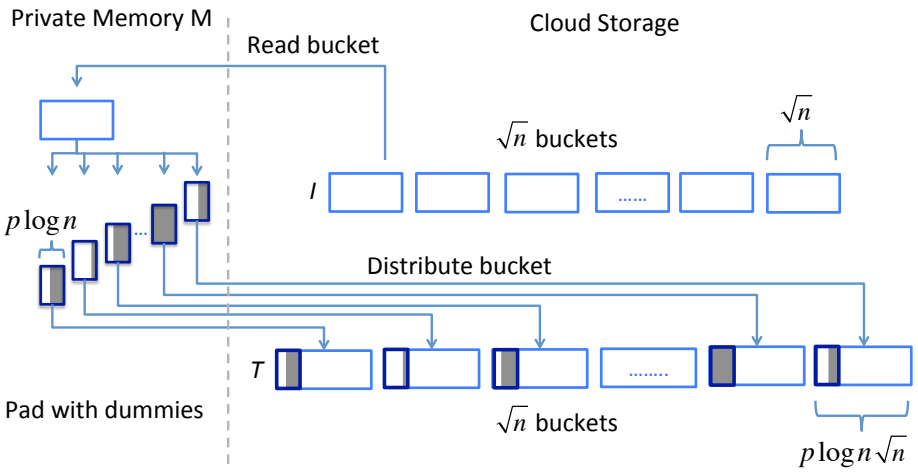

Figure 2: Illustration of the distribution phase of shuffle_pass (Algorithm 2). Shadowed regions represent dummy values added to pad each batch to the size of $p \log n$. The batches are encrypted, hence, one cannot tell where and how many dummy values there are in each batch.

in every bucket of T (rev_bucket$[\mathsf{id}_T]$ in the pseudo-code). Here, p is a constant and is determined in the analysis.

Each batch contains real and dummy elements. The first batch is filled in with real elements (x, v) that would go to the first bucket in O according to ρ, i.e., the elements for which $\lfloor \rho(x)/\sqrt{n} \rfloor = 0$. Similarly for every other batch. Since a bucket of I contains only $\sqrt{n}$ elements and we put $\sqrt{n} \times p \log n$ elements in total in all buckets in T, most batches will have less than $p \log n$ elements. We pad such batches with dummy elements to hide where and how many elements of I's bucket are placed in T (line 18). Note that a batch is re-encrypted before it is written to T, completely hiding the content and making it impossible to recognize where dummy or real elements are (lines 10 and 18). If according to ρ more than $p \log n$ elements are mapped from a bucket of I to a bucket of T, the algorithm fails (line 15). We later consider what happens in case of a failure. We note that $\sqrt{n}$ calls to putRange in the loop in lines 13-21 is for the ease of explanation only. These calls can be substituted by a single call putRangeDist, putting $\sqrt{n} \times p \log n$ elements all at once. Hence, for every bucket read from I, there is only *one* corresponding write to T.

Clean-up Phase The distribution phase leaves T with two problems: first, though the elements are in correct buckets according to ρ they are not in the correct locations inside the buckets, and second, T contains dummy elements. To remedy these problems, the clean-up phase in Algorithm 2, illustrated in Figure 3, proceeds by reading buckets of T of size $\sqrt{n} \times p \log n$ and writing in their place buckets of size $\sqrt{n}$.

Algorithm 2 Single pass $\mathsf{shuffle_pass}(I, T, \rho, O)$ of the Melbourne shuffle algorithm, where the user can read and store in private memory M upto $\sqrt{n} \times p \log n$ elements.

$\boldsymbol{I}$: array of n encrypted elements (x, v); $\boldsymbol{T}$: auxiliary array that fits $n \times p \log n$ encrypted elements, where p is a constant that is a parameter of the algorithm; $\boldsymbol{\rho}$: permutation; $\boldsymbol{O}$: permutation of I according to ρ, where every element is re-encrypted.

$\mathsf{max_elems} \leftarrow p \log n$
$\mathsf{num_buckets} \leftarrow \sqrt{n}$
{*Distribution phase: distribute elements of I into T*}
for $\mathsf{id}_I \in \{0, \ldots, \mathsf{num_buckets} - 1\}$ **do** {read buckets of I}
  $\mathsf{bucket}_M \leftarrow \mathsf{getRange}(I, \mathsf{id}_I \times \sqrt{n}, \sqrt{n})$
  $\mathsf{rev_bucket}_M \leftarrow \mathsf{empty_map}()$ {Reverse map of bucket ids in T to elements}
  for $e \in \mathsf{bucket}_M$ **do** {Assign elements their bucket ids in T}
    $(x, v) \leftarrow \mathsf{Dec}(e)$
    $\mathsf{id}_T \leftarrow \lfloor \rho(x) / \sqrt{n} \rfloor$ {Bucket id of element (x, v) in T according to its location in O}
    $\mathsf{rev_bucket}_M[\mathsf{id}_T].\mathsf{add}(\mathsf{Enc}(x, v))$ {Collect elements of same bucket}
  end for
  {Can be done via a single putRangeDist for $\sqrt{n}$ batches of size $\mathsf{max_elems}$}
  for $\mathsf{id}_T \in \{0, \ldots, \mathsf{num_buckets} - 1\}$ **do** {Distribute bucket_M in buckets of T}
    if $\mathsf{size}(\mathsf{rev_bucket}_M[\mathsf{id}_T]) > \mathsf{max_elems}$ **then**
      fail {ρ moves more than $p \log n$ elements from a bucket of I to a bucket of T}
    end if
    {Hide how many real elements go to T's buckets by padding with encrypted dummies}
    $\mathsf{rev_bucket}_M[\mathsf{id}_T] \leftarrow \mathsf{dummy_pad}(\mathsf{rev_bucket}_M[\mathsf{id}_T], \mathsf{max_elems})$
    {Write a batch of $\mathsf{max_elems}$ from every bucket of I to every bucket of T}
    $\mathsf{putRange}(T, \mathsf{id}_T \times \sqrt{n} \times \mathsf{max_elems} + \mathsf{max_elems} \times \mathsf{id}_I, \mathsf{rev_bucket}_M[\mathsf{id}_T])$
  end for
end for
{*Clean-up phase: clean T and write the result to O*}
for $\mathsf{id}_T \in \{0, \ldots, \mathsf{num_buckets} - 1\}$ **do** {read buckets of T}
  $\mathsf{bucket}_M \leftarrow \mathsf{getRange}(T, \mathsf{id}_T \times \sqrt{n} \times \mathsf{max_elems}, \sqrt{n} \times \mathsf{max_elems})$
  {Decrypt the bucket, remove dummy, sort real elements using ρ and re-encrypt}
  $\mathsf{bucket}_M \leftarrow \mathsf{clean}(\mathsf{bucket}_M)$
  {The distribution phase guarantees that bucket_M contains exactly $\sqrt{n}$ elements}
  $\mathsf{putRange}(O, \mathsf{id}_T \times \sqrt{n}, \mathsf{bucket}_M)$
end for

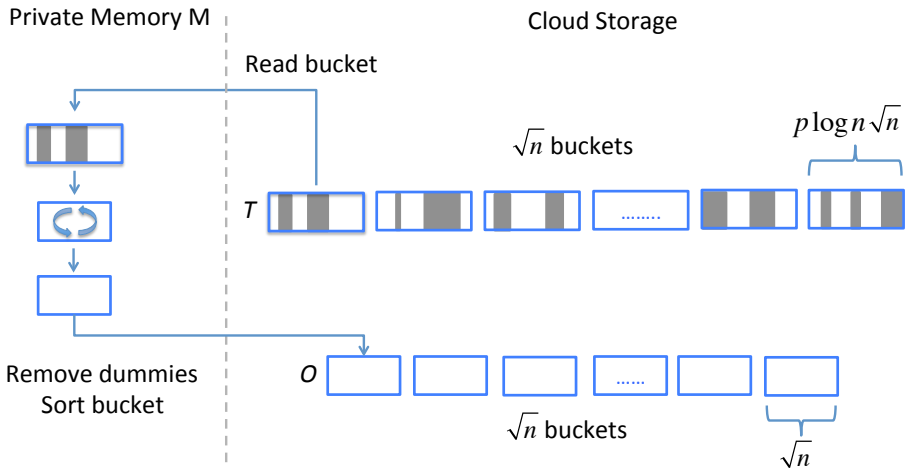

Figure 3: An illustration of the clean-up phase of shuffle_pass (Algorithm 2). Shadowed regions represent dummy values that are removed during the clean-up phase.

When processing each bucket, the algorithm removes dummy elements, sorts the remaining content of every bucket according to their final location in O (line 27). It is important to note that each written bucket contains exactly $\sqrt{n}$ elements before it is being written back. This follows from the fact that elements were distributed to buckets according to the permutation ρ and the algorithm failed in the distribution phase for those ρ that would have resulted in more than $\sqrt{n}$ elements in each bucket.

Performance The performance of the Melbourne shuffle is summarized in the following theorem.

Theorem 4.1: *Given an input array of size n, the Melbourne shuffle (Algorithm 1) executes $O(\sqrt{n})$ operations, each exchanging a message of size $O(\sqrt{n} \log n)$, between a user with private memory of size $O(\sqrt{n} \log n)$ and a server with storage of size $O(n \log n)$. Also, the user and server perform $O(n \log n)$ work.*

Proof: We first note that $\sqrt{n}$ calls to putRange in the loop in lines 13-21 is for the ease of explanation only. These calls can be substituted by a single call putRangeDist, putting $\sqrt{n} \times p \log n$ elements all at once.

A single shuffle pass in Algorithm 2 requires $2\sqrt{n}$ calls to getRange, $\sqrt{n}$ calls to putRangeDist, and $\sqrt{n}$ calls to putRange, assuming the user and the server can exchange up to $\sqrt{n} \times p \log n$ elements in a single request. The shuffle in Algorithm 1 requires $8\sqrt{n}$ requests in total since it makes 2 calls to the shuffle pass procedure. The private memory required at the user to perform the shuffle is $\sqrt{n} \times p \log n$. The required server's memory is $n \times p \log n$. However, this overhead is temporary since the increase in memory happens only during the shuffle pass and is reduced to n when the shuffle is finished. Similarly for the user, the memory of size $\sqrt{n} \times p \log n$ is required only during the shuffle. We note that the total computation for the user and the server is $O(n \log n)$. ■

4.3 Security Analysis

In this section, we show that the Melbourne shuffle (Algorithm 1) is oblivious for every permutation π with high probability.

Definition 4.2: *Let A be an array of n elements such that every $x \in [1, n]$ is at location $\pi_0(x)$ in A. Let B be an array that stores a permutation π of elements in A, i.e., $B = \pi(A)$. Split A and B in $\sqrt{n}$ buckets of equal size and fix a constant $p \geq \mathrm{e}$. Let π be a permutation on n elements where every bucket of B contains at most $p \log n$ elements of every bucket of A. We refer to the set of all such permutations as $P(\pi_0)$.*

Lemma 4.3: *The size of set $P(\pi_0)$ is $(1 - \mathsf{negl}(n)) \times n!$, for every permutation π_0.*

Proof: Let π be a random permutation from all possible $n!$ permutations. We consider the relationship between the input array A and a permutation of A, $B = \pi(A)$. We start by splitting A and B in buckets of size $\sqrt{n}$ and numbering buckets from 1 to $\sqrt{n}$ using their order in each array. The analysis below estimates how many permutations can be constructed by restricting the maximum number of elements from a bucket of A appearing in any bucket of B to $p \log n$.

Let X_a^b be a random variable that measures the number of elements from ath bucket of A present in bth bucket of B. The mean value of X_a^b is 1, since we are distributing $\sqrt{n}$ elements of a among $\sqrt{n}$ buckets of B. Although X_a^b, for $1 \leq a, b \leq \sqrt{n}$, variables are dependent between each other, we can use the Poisson Approximation [19, Chapter 5.4] and instead work with n independent Poisson random variables Y_a^b with mean 1.

Given n variables Y_a^b we are interested in bounding the probability of the event that there is no a and b such that $Y_a^b \geq p \log n$. For a specific a and b it is:

$$\Pr[Y_a^b \geq p \log n] \leq \frac{1}{\mathrm{e}} \left(\frac{\mathrm{e}}{p \log n} \right)^{p \log n}.$$

Using union bound, the probability that at least one of the Y_a^bs is greater than $p \log n$ is at most

$$n \frac{1}{\mathrm{e}} \left(\frac{\mathrm{e}}{p \log n} \right)^{p \log n}.$$

Since Y_a^bs are a Poisson approximation of the variables X_a^b, the probability that at least one of the X_a^bs is greater than $p \log n$ is at most

$$2n \frac{1}{\mathrm{e}} \left(\frac{\mathrm{e}}{p \log n} \right)^{p \log n}.$$

Setting $p \geq \mathrm{e}$ we get

$$2n \frac{1}{\mathrm{e}} \left(\frac{\mathrm{e}}{p \log n} \right)^{p \log n} \leq \frac{2n}{(\log n)^{p \log n}} = \frac{2n}{n^{p \log \log n}} = 2\mathsf{negl}(n) = \mathsf{negl}(n).$$

■

Lemma 4.4: *Let π_0 be the initial permutation of n elements in the input array I. Method* shuffle_pass *(Algorithm 2) succeeds for all permutations $\rho \in P(\pi_0)$.*

Proof: The algorithm allocates elements of I in O according to ρ by first putting them into corrects buckets (lines 4–22) and then sorting every bucket using ρ (line 27). By construction the algorithm fails for any permutation ρ that requires more than $p \log n$ elements from a bucket of I mapped to buckets of O (line 14–16). ■

Lemma 4.5: *Method* $\mathsf{shuffle}(I, \pi, O)$ *(Algorithm 1) is a randomized shuffle algorithm that succeeds with very high probability.*

Proof: Let π_0 be the initial permutation of the input array I. Algorithm 1 makes two calls to shuffle pass which succeeds for all possible permutations except for $1/n^{\Omega(\log\log n)}$ fraction of them (Lemmas 4.3 and 4.4). The first shuffle pass is executed for permutation π_1 on an input permuted according to some permutation π_0. Since π_1 is picked using internal random coins, the probability of the shuffle pass failing is independent of π_0 and is bounded by $1/n^{\Omega(\log\log n)}$. If the first shuffle pass did not fail, the shuffle pass is executed second time with input permutation π. The second shuffle is executed on the input array that is permuted according to a random permutation $\pi_1 \in P(\pi_0)$. The second pass does not fail iff $\pi_1 \in P(\pi)$. Hence, Algorithm 1 fails if $\pi_1 \notin P(\pi_0)$ or $\pi_1 \notin P(\pi)$. By Lemma 4.3, the probability of either of these events is negligible in n, hence the basic Melbourne shuffle succeeds with very high probability for any π_0 and π. ■

We show that method shuffle_pass (Algorithm 2) is oblivious by mapping it to the Oblivious Shuffle Model in Section 3.1, extracting the corresponding transcript and showing that the transcript reveals no information about the underlying permutation if the encryption scheme is CPA secure (see Section 2.1).

Method shuffle_pass (Algorithm 2) corresponds to GenRequest in the model and calls to getRange and putRange trigger calls to GenResponse at the server. We do not describe GenResponse since it depends on the implementation details of the remote storage provider. We are only interested in the fact that it uses server's state S to store and maintain arrays I, T and O. The transcript α of the shuffle execution is defined as follows. The request r_i is either $\mathsf{getRange}(S, i, l)$ or $\mathsf{putRange}(S, i, a)$, e.g., in line 29 (S, x, a) is $(O, \mathsf{id}_T \times \sqrt{n}, \mathsf{bucket}_M)$. The response g_i to getRange is an array a, e.g., a contains $\sqrt{n}$ elements in line 5 and is stored in bucket_M. The response to putRange is empty. We first analyze the metadata (Definition 2.1) that corresponds to every request between the client and the server, and show that, unless the algorithm fails, they depend on the size of the input only, and are independent from the input array and the desired permutation. Hence, we obtain that the Melbourne shuffle is a data independent shuffle algorithm. We finally show that if the content exchanged is encrypted, as it is in method shuffle_pass (Algorithm 2), the Melbourne shuffle (Algorithm 1) is oblivious.

Lemma 4.6: *The metadata of requests exchanged between the client and the server in method* shuffle_pass *(Algorithm 2) is independent of permutation π_0 of the input I and output permutation $\rho \in P(\pi_0)$, and depends only on n.*

Proof: Let α be a sequence of (request, response) pairs exchanged between the client and the server, denoted as (r_i, g_i), where r_i is either a getRange or putRange and g_i is bucket_M for getRange and empty for getRange. The sequence α can be further split in $\sqrt{n}$ (getRange, putRange) calls that correspond to distribution phase and $\sqrt{n}$ (getRange, putRange) during the clean-up phase.

The metadata of putRange is the name of the array, the location within the array and how many elements should be read. In the algorithm these correspond to reading an array of size n sequentially in buckets of size $\sqrt{n}$ (distribution phase, line 5) and $\sqrt{n} \times p \log n$ (clean-up phase, line 25). These data depend only on n and p. The metadata for a putRange call consists of the array to be accessed, the location where to put the data and the size of the data to be written. In the algorithm, first calls to putRange place $\sqrt{n}$ batches of size $p \log n$ in the temporary array to locations that depend on the input bucket that has been read using getRange (line 20). These locations are deterministic since getRange simply scans the input array. The second sequence of calls to putRange happens during the clean-up phase when buckets of size $\sqrt{n}$ are written sequentially to the output array (line 29). These calls are also deterministic. It is also easy to show that the transcript where a sequence of $\sqrt{n}$ calls to putRange in lines 13-21 is substituted with a single call to putRangeDist is also deterministic. ■

Lemma 4.7: *Let π_0 be the initial permutation of n elements in the input array I and let ρ be a permutation from the set $P(\pi_0)$. Method* shuffle_pass(I, T, ρ, O) *(Algorithm 2) is an oblivious shuffle according to Definition 3.3.*

Proof: We showed in Lemma 4.4 that Algorithm 2 succeeds for all $\rho \in P(\pi_0)$. We also showed that all metadata in the transcript that is revealed to the adversary $\mathcal{A}$ in the Shuffle-IND game in Definition 3.2 after every call to Shuffle is independent of data content and hence can be determined based only on n and is the same for any choice of input and output. The data content exchanged in each call does depend on the data, however, it is always encrypted. We show that the security of the shuffle depends on the security of the underlying encryption scheme.

To the contrary, we assume that there is a PPT adversary $\mathcal{A}$ that can distinguish with a non-negligible advantage two permutations τ_0 and τ_1 by observing the transcript of one of them. Hence, this adversary can win Shuffle-IND game. We show that if $\mathcal{A}$ exists then we can construct an adversary $\mathcal{B}$ who can use $\mathcal{A}$ to win Enc-IND-CPA game with a non-negligible advantage, which would break our assumption about the encryption scheme. We recall that in Enc-IND-CPA game $\mathcal{B}$ has access to oracles Enc and Dec and can encrypt and decrypt sequences of messages of his choice, except asking for the decryption of the challenge ciphertext.

We construct the adversary $\mathcal{B}$ as follows. $\mathcal{B}$ does not need to run Setup since he has oracle access to Enc and Dec. $\mathcal{A}$ starts making calls to Shuffle on chosen pairs of input arrays and permutations. $\mathcal{B}$ imitates the shuffle by responding with the encrypted permutation and a transcript α, that he can produce himself. He continues doing so until $\mathcal{A}$ comes up with a challenge of two (input, permutation) pairs (A_0, τ_0) and (A_1, τ_1). $\mathcal{B}$ first creates a static transcript that will be the same for both permutations. He then, extracts all the calls to be made to Enc into two sequences: one that corresponds to (A_0, τ_0) and one to (A_1, τ_1). He gives these two sequences of elements to be encrypted to his own challenger. The challenger picks one sequence at random, encrypts all its plaintexts (i.e., elements) one by one and gives the result to $\mathcal{B}$. $\mathcal{B}$ combines the ciphertexts with the metadata, that depend only on n, to create a valid transcript α and sends it to $\mathcal{A}$. He continues, responding to Shuffle requests from $\mathcal{A}$ until $\mathcal{A}$ outputs his guess b for the pair (A_b, τ_b). $\mathcal{B}$ outputs b as his guess for which sequence of messages his challenger picked. Since $\mathcal{A}$'s advantage in winning the game is non-negligible so is $\mathcal{B}$'s. ■

Theorem 4.8: *The Melbourne Shuffle (Algorithm 1) is a randomized shuffle algorithm that succeeds with very high probability and is data-oblivious according to Definition 3.3.*

Proof: Algorithm 1 makes two calls to method shuffle_pass, which is oblivious by Lemma 4.7. If Algorithm 1 is not oblivious, then there is a PPT adversary $\mathcal{A}$ who can distinguish shuffle of two permutations. If $\mathcal{A}$ exists, we can build an adversary $\mathcal{B}$ who can break the security of the underlying shuffle_pass using $\mathcal{A}$. Whenever $\mathcal{A}$ makes a call to Shuffle for a permutation π, $\mathcal{B}$ first picks a random permutation π_1 and calls shuffle_pass on π_1. It then uses the output of this call and π to make another call to shuffle_pass, this time returning the output to $\mathcal{A}$. This continues until $\mathcal{A}$ comes up with two challenge pairs (input, permutation) (A_0, τ_0) and (A_1, τ_1). $\mathcal{B}$ first picks a random permutation and runs shuffle-pass on it. He then gives his own challenger the output of this shuffle along with (A_0, τ_0) and (A_1, τ_1). The challenger returns to $\mathcal{B}$ the shuffle according to (A_b, τ_b), keeping b secret. $\mathcal{B}$ forwards what he receives to $\mathcal{A}$. $\mathcal{B}$ continues replying shuffle requests of $\mathcal{A}$ as he did before until $\mathcal{A}$ makes a guess for b. $\mathcal{B}$ outputs this guess as his own. ■

Note that method shuffle_pass outputs fail for some permutations. Whenever it does so, $\mathcal{B}$ sends this information to $\mathcal{A}$. However, as showed in Lemma 4.5 this happens with negligible probability for any pair of input and output permutations and reveals nothing to the adversary about the output permutation since the failure is due to secret random bits.

5 The Optimized Melbourne Shuffle

In this section we present an optimized version of the Melbourne shuffle that has smaller memory requirements on the memory of the user and the server, and succeeds with higher probability than the basic version of the previous section. The framework of the optimized version is similar: we first re-randomize the input, i.e., shuffle it according to a random permutation π' and then shuffle π' towards the desired permutation π. The main difference with the basic version lies in the shuffle pass.

5.1 Algorithm

As in the basic version, the shuffle pass splits the input array I and the output array O in consequent buckets of size $\sqrt{n}$. For auxiliary storage we use two temporary arrays T_1 and T_2 of size $p_1 n$ and $p_2 n$, respectively, where $p_1, p_2 > 1$ are constants to be determined in the analysis. We split T_1 and T_2 in buckets of size $p_1\sqrt{n}$ and $p_2\sqrt{n}$, respectively. The shuffle pass proceeds with two distribution phases, instead of one for the basic version, followed by a single clean-up phase in the end. Detailed pseudo-code is given in Algorithms 3–6

Algorithm 3 The optimized Melbourne shuffle algorithm, optim_shuffle(I, π, O), where the user can read and store in private memory M up to $p_2\sqrt{n}$ elements.

$\boldsymbol{I}$: an array of n encrypted elements (x, v); $\boldsymbol{\pi}$: the desired permutation for elements of I; $\boldsymbol{O}$: the output array containing the permutation of I according to π, where every element (x, v) is re-encrypted.

1: Let π_1 be a random permutation
2: Let T_1 and T_2 be two empty arrays stored remotely of size $p_1 n$ and $p_2 n$, respectively
3: optim_shuffle_pass(I, T_1, T_2, π_1, O)
4: $I \leftarrow O$
5: optim_shuffle_pass(I, T_1, T_2, π, O)

Algorithm 4 Single pass optim_shuffle_pass(I, T_1, T_2, ρ, O) of the optimized Melbourne shuffle algorithm, where the user can read and store in private memory M up to $p_2\sqrt{n}$ elements.

$\boldsymbol{I}$: an array of n encrypted elements (x, v); $\boldsymbol{T_1}$: an auxiliary array that fits $p_1 n$ elements; $\boldsymbol{T_2}$: an auxiliary array that fits $p_2 n$ elements; $\boldsymbol{\rho}$: the desired permutation for elements of I; $\boldsymbol{O}$: the output array containing the permutation of I according to ρ, where every elements (x, v) is re-encrypted.

1: distr_phase1(I, ρ, T_1)
2: distr_phase2(T_1, ρ, T_2)
3: clean_up_phase(T_2, ρ, O)

The first distribution phase moves elements from I to T_1, the second distribution phase moves elements from T_1 to T_2. We abstract the layout of elements in each array further by sequentially splitting buckets in chunks. The goal of the first distribution phase is to place elements in correct chunks and place them in correct buckets within the chunks in the second distribution phase. When a bucket is read, it is decrypted and any elements that are written back are re-encrypted. In the following, we denote with ρ the target distribution of the shuffle pass. For an array of size n, this algorithm assumes messages and client private memory of size $O(\sqrt{n})$ and server memory of size $O(n)$.

Distribution Phase I We view a sequence of $\sqrt[4]{n}$ buckets in each array as a chunk. Hence, I, T_1, T_2 and O each have $\sqrt[4]{n}$ chunks. The goal of the first distribution phase is to place elements of I in T_1 in such a

way that all elements that belong to the first chunk of O according to ρ can be found in the first chunk of T_1, similarly for the second chunk, and so on. The pseudo-code for this phase is given in Algorithm 5.

Referring to the pseudo-code of Algorithm 5, the distribution proceeds by placing elements in chunks at the bucket level (line 7). For every bucket in I, a batch with at most $p_1\sqrt[4]{n}$ elements is moved to every chunk of T_1. The assignment of elements to chunks of T_1 is determined by the permutation ρ that we wish to achieve. In particular, an element (x, v) is placed in the $\lfloor \rho(x)/n^{3/4} \rfloor$th chunk of T_1 (line 11). Note that $\lfloor \rho(x)/n^{3/4} \rfloor$ is the index of the chunk where (x, v) belongs in O since $n^{3/4}$ is the size of a chunk in O.

As a result of placing $\sqrt{n}$ elements in $\sqrt[4]{n}$ batches of size $p_1\sqrt[4]{n}$, some batches will not be full. We pad such batches with dummy elements (line 21). If, on the other hand, more than $p_1\sqrt[4]{n}$ elements need to be moved according to ρ, the algorithm fails. We consider and analyze this case later. If the distribution succeeds, every chunk of T_1 has exactly $p_1 n^{3/4}$ elements with $n^{3/4}$ real elements only. Note that here, as in the distribution phase of the basic shuffle, $\sqrt[4]{n}$ calls in lines 15-24 can be substituted with a single putRangeDist of size $p_1\sqrt{n}$. For an illustration of this phase refer to Figure 4.

Algorithm 5 distr_phase1(I, ρ, T_1) of the optimized Melbourne shuffle algorithm, where the user can read and store in private memory M upto $p_2\sqrt{n}$ elements, $p_1 \leq p_2$.

$\boldsymbol{I}$: an array of n encrypted elements (x, v); $\boldsymbol{\rho}$: the desired permutation; $\boldsymbol{T_1}$: output array containing encrypted elements in correct chunks according to ρ.

1: max_elems $\leftarrow p_1\sqrt[4]{n}$
2: num_chunks $\leftarrow \sqrt[4]{n}$
3: num_buckets $\leftarrow \sqrt{n}$
4: chunk_size $\leftarrow n^{3/4}$
5: {Distribution phase I: distribute elements of I in T_1 according to their chunk index in O}
6: **for** $i \in \{0, \ldots,$ num_buckets $- 1\}$ **do** {go through buckets in I}
7: bucket$_M \leftarrow$ getRange$(I, i \times \sqrt{n}, \sqrt{n})$
8: rev_bucket$_M \leftarrow$ empty_map() {Reverse map of chunk ids in T_1 to elements}
9: **for** $e \in$ bucket$_M$ **do** {Assign elements their chunk ids according to ρ}
10: $(x, v) \leftarrow$ Dec(e)
11: cid $\leftarrow \lfloor \rho(x)/$chunk_size$\rfloor$ {Chunk id of element (x, v) in the output shuffle}
12: rev_bucket$_M$[cid].add(Enc(x, v)) {Collect elements that correspond to the same chunk in T_1}
13: **end for**
14: {Can be done via a single putRangeDist for $\sqrt[4]{n}$ batches of size max_elems}
15: **for** cid $\in \{0, \ldots,$ num_chunks $- 1\}$ **do** {Distribute bucket$_M$ among chunks of T_1}
16: **if** size(rev_bucket$_M$[cid]) $>$ max_elems **then**
17: {Permutation requires more than $p_1\sqrt[4]{n}$ elements moved from a bucket of I to a chunk of T_1}
18: **fail**
19: **end if**
20: {Hide how many real elements go to T_1 from this bucket by padding with encrypted dummies}
21: rev_bucket$_M$[cid] $\leftarrow$ dummy_pad(rev_bucket$_M$[cid], max_elems)
22: {Write a batch of max_elems from every bucket of I to every chunk of T_1}
23: putRange$(T_1,$ cid $\times$ chunk_size $+ i \times$ max_elems, rev_bucket$_M$[cid]$)$
24: **end for**
25: **end for**

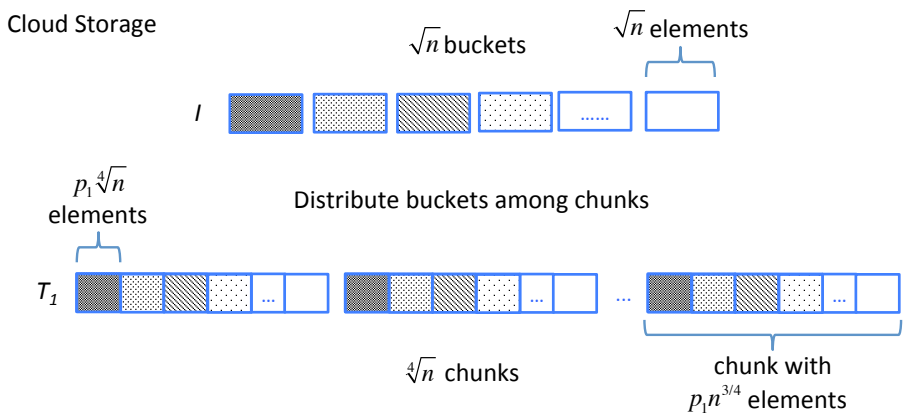

Figure 4: Illustration of the arrangement of elements from the input I in the output T_1 after the first distribution phase of the Optimized Melbourne Shuffle (Section 5). See pseudocode in Algorithm 5.

Distribution Phase II Observe that elements in T_1 belong to the correct chunk but not the correct bucket within the chunk. The second distribution phase remedies this, such that by the end of this phase elements of chunks of T_1 are in their correct buckets in T_2. The pseudo-code of this phase is presented in Algorithm 6.

Referring to the pseudo-code of Algorithm 6, we proceed by reading buckets of size $p_1\sqrt{n}$ in each chunk of T_1 (line 10). For every jth bucket of ith chunk of T_1 we write $\sqrt[4]{n}$ batches of size $p_2\sqrt[4]{n}$ to T_2 (line 28). The batches are written to all $\sqrt[4]{n}$ buckets of ith chunk in T_2. A real element (x, v) is assigned to $\lfloor \rho(x)/\sqrt{n} \rfloor \mod n^{3/4}$th batch (line 15), which is the bucket id in the ith chuck of T_2, where this batch will be written to (line 28). Note that dummy elements added during the first distribution phase are ignored. If there is a batch that no element has been assigned to or if a batch has less than $p_2\sqrt[4]{n}$ elements, we pad it with dummies. If there is a batch with more than $p_2\sqrt[4]{n}$ elements, the algorithm fails. Note that $i \times \sqrt[4]{n} + \text{bid}$ is the index of x's bucket in O.

Once all the batches are set up, the batches are written in their place within the same chunk in T_2. Again, $\sqrt[4]{n}$ calls in lines 20-29 can be substituted with a single putRangeDist of size $p_2\sqrt{n}$. When all buckets of T_1 have been processed, buckets of T_2 contain $p_2\sqrt{n}$ elements each. This is a consequence of writing a batch of $p_2\sqrt[4]{n}$ elements from every chunk of T_1. Moreover, every bucket of $p_2\sqrt{n}$ elements contains all $\sqrt{n}$ elements that belong to this bucket in O. For an illustration of this phase refer to Figure 5.

Clean-up Phase This phase is similar to the clean-up phase of the basic version, hence we omit its pseudo-code. Recall that a bucket of T_2 contains $p_2\sqrt{n}$ encrypted elements but contains only $\sqrt{n}$ real elements which are in the correct bucket but not in the correct spot within the bucket. We remedy this by reading every bucket $j \in [1, \sqrt{n}]$, decrypting it, removing dummy elements such that only $\sqrt{n}$ real elements are left, sorting it according to ρ, re-encrypting the elements and writing them back to the jth bucket of O.

Performance The performance of the optimized Melbourne shuffle is summarized in the following theorem.

Theorem 5.1: *Given an input array of size n, the optimized Melbourne shuffle (Algorithm 3) executes $O(\sqrt{n})$ operations, each exchanging a message of size $O(\sqrt{n})$ between a user with private memory of size $O(\sqrt{n})$ and a server with storage of size $O(n)$. Also, the user and the server perform $O(n)$ work.*

Algorithm 6 distr_phase2(T_1, ρ, T_2) of the optimized Melbourne shuffle algorithm, where the user can read and store in private memory M upto $p_2\sqrt{n}$ elements.

$\boldsymbol{T_1}$: array containing encrypted elements in correct chunks according to ρ; $\boldsymbol{\rho}$: the desired permutation; $\boldsymbol{T_2}$: output array containing encrypted elements in correct buckets according to ρ.

max_elems $\leftarrow p_2 \sqrt[4]{n}$
num_chunks $\leftarrow \sqrt[4]{n}$
num_buckets $\leftarrow \sqrt{n}$
bucket_size $\leftarrow \sqrt{n}$
chunk_size $\leftarrow n^{3/4}$
num_buckets_per_chunk $\leftarrow \sqrt[4]{n}$
{Distribution phase II: distribute elements of T_1 in T_2}
for $i \in \{0, \ldots,$ num_chunks $- 1\}$ **do** {go through chunks in T_1}
  for $j \in \{0, \ldots,$ num_buckets_per_chunk $- 1\}$ **do** {go through buckets in ith chunk}
    bucket$_M \leftarrow$ getRange$(I, i \times$ chunk_size $+ j \times \sqrt{n}, p_1\sqrt{n})$
    rev_bucket$_M \leftarrow$ empty_map() {Reverse map of bucket ids in T_2 to elements}
    for $e \in$ bucket$_M$ **do** {Assign elements their bucket ids according to ρ}
      $(x, v) \leftarrow$ Dec(e)
      if (x, v) is real **then** {Ignore dummy elements}
        bid $\leftarrow \lfloor \rho(x)/$bucket_size$\rfloor \mod$ num_buckets_per_chunk {Bucket id of element (x, v) in ith chunk of T_2}
        rev_bucket$_M$[bid].add(Enc(x, v)) {Collect elements that correspond to the same bucket}
      end if
    end for
    {Can be done via a single putRangeDist for $\sqrt[4]{n}$ batches of size max_elems}
    for bid $\in \{0, \ldots,$ num_buckets_per_chunk $- 1\}$ **do** {Distribute bucket$_M$ among buckets of T_2 in ith chunk)}
      if size(rev_bucket$_M$[bid]) $>$ max_elems **then**
        {Permutation requires more than $p_2 \sqrt[4]{n}$ elements moved from T_1 to T_2}
        fail
      end if
      {Hide how many real elements go to T_2 from this bucket by padding with encrypted dummies}
      rev_bucket$_M$[bid] $\leftarrow$ dummy_pad(rev_bucket$_M$[bid], max_elems)
      {Write a batch of max_elems from every bucket of T_1 to buckets in ith chunk in T_2}
      putRange$(T_2, i \times$ chunk_size $+$ bid $\times p_2\sqrt{n} + j \times$ max_elems, rev_bucket$_M$[bid]$)$
    end for
  end for
end for

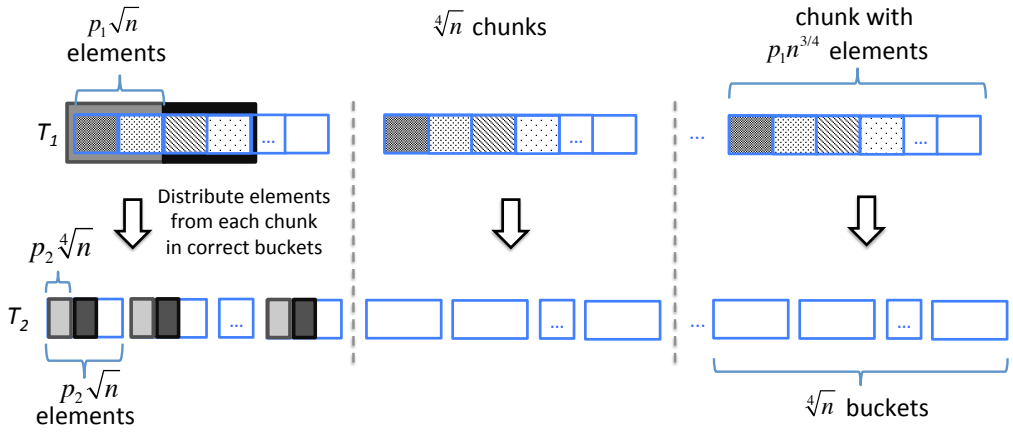

Figure 5: Illustration of the arrangement of elements after processing two buckets of the first chunk of the input T_1 in the output T_2 in the second distribution phase of the optimized Melbourne Shuffle. See pseudocode in Algorithm 6. Read and written elements are indicated in gray and black colors.

Proof: A single optimized shuffle pass in Algorithm 4 requires $3\sqrt{n}$ calls to getRange, $2\sqrt{n}$ calls to putRangeDist and $\sqrt{n}$ calls to putRange when the user and the server can exchange up to $p_2\sqrt{n}$ elements in a single request. Then, the optimized shuffle algorithm in Algorithm 3 requires 12 requests in total. The size of the user private memory required to perform the shuffle is at most $p_2\sqrt{n}$, while required server's memory is at most $(p_1 + p_2)n$. However, this overhead in server's memory is temporary since this memory is required only during the the distribution phases. Similarly, for the user private memory, the size of $O(\sqrt{n})$ is required while shuffling and is decreased to $O(1)$ once it is finished. The total work for the user is $O(n)$. ■

5.2 Security Analysis

In this section, we show that the optimized Melbourne shuffle is oblivious for every permutation π with very high probability.

Definition 5.2: *Let A be an array of n elements such that for every $1 \leq x \leq n$ is at location $\pi_0(x)$ in A. Let $B = \pi(A)$ be a permutation on A. Split A and B in $\sqrt[4]{n}$ chunks of equal size and fix constants $p_1, p_2 \geq 1$. Split every chunk of A and B further in $\sqrt[4]{n}$ buckets of size $\sqrt{n}$.*

We say that a permutation $\pi \in Q(\pi_0)$ if according to π: (1) every chunk of B has at most $p_1\sqrt[4]{n}$ elements from every bucket of A, and (2) if every bucket of B has at most $p_2\sqrt[4]{n}$ elements from every chunk of A.

Lemma 5.3: *The size of $Q(\pi_0)$ is $(1 - \mathsf{negl}(n)) \times n!$ for every permutation π_0.*

Proof: Let A and B be arrays of n items. Split A and B in chunks of size $n^{3/4}$ and then split each chunk in $\sqrt[4]{n}$ buckets of size $\sqrt{n}$. The shuffle pass succeeds for any permutation that satisfies the following two constraints. (1) Every chunk of B contains at most $p_1\sqrt[4]{n}$ elements from every bucket of A and (2) Each bucket of B has at most $p_2\sqrt[4]{n}$ elements of the ith chunk of A. We compute the probability of a uniform permutation not satisfying each of the constrains independently and use union bound to bound the probability of failing at least one of the constraints.

Let X_a^b be the number of elements of ath bucket of A in bth chunk of B. Since there are $\sqrt{n}$ elements in every bucket of A and there are $\sqrt[4]{n}$ chunks in B, the mean of X_a^b is $\sqrt[4]{n}$. As in Lemma 4.3 we use

the Poisson approximation and use independent Poisson variables Y_a^b with mean value of $\sqrt[4]{n}$. Then for a specific a and b:

$$\Pr[Y_a^b \geq p_1 \sqrt[4]{n}] \leq \frac{\mathrm{e}^{-\sqrt[4]{n}}(\mathrm{e}\sqrt[4]{n})^{p_1\sqrt[4]{n}}}{(p_1\sqrt[4]{n})^{p_1\sqrt[4]{n}}} = \frac{\mathrm{e}^{p_1}\sqrt[4]{n}^{p_1\sqrt[4]{n}}}{(p_1\sqrt[4]{n})^{p_1\sqrt[4]{n}}} = \frac{\mathrm{e}^{p_1}}{{p_1}^{p_1\sqrt[4]{n}}}.$$

Using union bound the probability:

$$\Pr[\bigcup_{\substack{1 \leq a \leq \sqrt{n} \\ 1 \leq b \leq \sqrt[4]{n}}} Y_a^b \geq p_1\sqrt[4]{n}] \leq n^{3/4}\frac{\mathrm{e}^{p_1}}{{p_1}^{p_1\sqrt[4]{n}}} = \mathsf{negl}[n].$$

Hence,

$$\Pr[\bigcup X_a^b \geq p_1\sqrt[4]{n}] \leq 2\Pr[\bigcup Y_a^b \geq p_1\sqrt[4]{n}] = \mathsf{negl}[n].$$

Consider the second constraint. Let U_c^d be the the number of elements of cth chunk of A in dth bucket of B. Since there are $n^{3/4}$ elements in every chunk of A and B has $\sqrt{n}$ buckets, U_c^d is a Poisson random variable with parameter $\sqrt[4]{n}$. Using the Poisson approximation, let W_c^d be $n^{3/4}$ independent Poisson random variables with mean $\sqrt[4]{n}$. Similar to the analysis of the first constraint, the probability that there is at least one bucket of B that has more than $p_2\sqrt[4]{n}$ elements of a chunk of A is $\mathsf{negl}(n)$.

Using union bound the probability that at least one of the constraints is not satisfied is the sum of the respective probabilities, hence, with probability $1 - \mathsf{neg}(n)$, moving elements for a random permutations via two levels will succeed. ■

Lemma 5.4: *The metadata of requests exchanged between the client and the server in method* optim_shuffle_pass *(Algorithm 4) is independent of input permutation π_0 and output permutation $\rho \in Q(\pi_0)$, and depends only on n.*

Proof:[Sketch] We consider the metadata from getRange and putRange requests during the two distribution phases and the clean-up phase. During the distribution phase I $\sqrt{n}$ getRange calls are made with metadata representing sequential get accesses to blocks of size $\sqrt{n}$ of input I (line 7 in Algorithm 5). The call to putRangeDist puts blocks of data of size $p_1\sqrt{n}$ to locations in T_1 that are determined based on the id of the block being read from I (lines 15–24 in Algorithm 5). Hence, metadata depend on the size of I, n, and does not depend on the content.

During the distribution phase II, blocks of size $p_1\sqrt{n}$ are sequentially read using getRange on T_1 (line 10 in Algorithm 6) and block are written back using putRangeDist to T_2 of size $p_2\sqrt{n}$ (lines 20–29 in Algorithm 6). Hence, metadata is deterministic based on n.

Clean-up phase sequentially reads blocks of size $p_2\sqrt{n}$ from T_2 and writes back blocks of size $\sqrt{n}$, independent of the data. Hence, metadata, produced by Algorithm 4 is data independent. ■

Lemma 5.5: *Let π_0 be the initial permutation of n elements in the input array I. Algorithm 4 succeeds for all permutations $\rho \in Q(\pi_0)$ and is data-oblivious according to Definition 3.3.*

Proof:[Sketch] The algorithm can fail during the first or second distribution phases (lines 18 in Algorithm 5 and 23 in Algorithm 6). The first failure happens when distributing a bucket of size $\sqrt{n}$ of I among $\sqrt[4]{n}$ buckets in T_1 and there is a bucket that needs to put more than $p_1\sqrt{n}$ of its elements to one of the buckets according to the input permutation ρ. This failure represents failure to satisfy the first condition of Definition 5.2 of $Q(\pi_0)$.

The second distribution phase results in a failure when one of the buckets in T_1, formed in the previous step, is distributed among $\sqrt[4]{n}$ buckets of T_2 and requires more than $p_2\sqrt[4]{n}$ of its elements in one of the buckets of T_2. This corresponds to failing the second condition of Definition 5.2. Hence, Algorithm succeeds for all permutations in $Q(\pi_0)$.

The metadata produced by the algorithm is data-independent as showed in Lemma 5.4. Since data is re-encrypted every time it is written back in getRange we can use the reduction similar to Lemma 4.7 to show that the security of the shuffle depends on Enc-IND-CPA secure encryption scheme. ■

Theorem 5.6: *The optimized Melbourne Shuffle (Algorithm 3) is a randomized shuffle algorithm that succeeds with very high probability and is data-oblivious according to Definition 3.3.*

Proof:[Sketch] Algorithm 3 runs the optimized shuffle pass algorithm in Algorithm 4 twice. Each shuffle pass succeeds if the desired permutation is from the set $Q(\pi_0)$, where π_0 is the original permutation of the input I. Running the shuffle twice: once for a random permutation π_1 and then for the desired permutation, ensures that the algorithm can succeed for every permutation with very high probability, where the probability depends on random coin tosses that determine the intermediate random permutation π_1. We showed in Lemma 5.5 that Algorithm 4 is data-oblivious, hence, running it twice also produces a data-oblivious algorithm (same reduction argument as in Theorem 4.8). ■

6 The Melbourne Shuffle with Small Messages

The Melbourne shuffle and its optimized version can be extended to work with messages and private memory of size $n^{1/c}\log n$ (or $n^{1/c}$ for the optimized version), for $c \geq 3$.

The idea behind the approach is to run the algorithm recursively with depth $c-1$. For a fixed c, one first splits the output in large buckets of size $n^{(c-1)/c}$ and executes the shuffle as in the square root case: distributing $n^{1/c}$ among $n^{1/c}$ large buckets. We call this the first level of the recursion. After this, each large bucket has correct elements but not in correct buckets nor positions. The square root shuffle is executed again on each large bucket, but now splitting the large bucket of size $O(n^{(c-1)/c}\log n)$ ($O(n^{(c-1)/c})$ for the optimized version) in $n^{(c-2)/c}$ buckets and again using only $n^{1/c}$ private memory. The client follows the recursion until the size of the inner buckets becomes $O(n^{2/c})$ when elements can be distributed in their correct buckets of size $n^{1/c}$. At this point, the buckets are small enough that they can be read to private memory during the clean-up phase and be placed in correct positions within their bucket. Hence, there are $c-1$ levels of recursion. Each level i requires a block of $n^{(c-i)/c}$ buckets, each of size $n^{1/c}$, to be distributed among $n^{1/c}$ output buckets. Since every level has $n^{(i-1)/c}$ blocks, the square root solution is required to be executed $n^{(i-1)/c}$ times per level, each making $O(n^{(c-i)/c})$ accesses. Hence, the total number of requests can be expressed as $\sum_{i=1}^{c-1} O(n^{(c-i)/c} \times n^{(i-1)/c}) = O((c-1)n^{(c-1)/c})$.

For the Melbourne shuffle that uses private memory and messages with a multiplicative $\log n$ factor (Section 4), the naive solution could accumulate the $O((\log n)^c)$ factor if used naively. This happens due to reading $n^{1/c}$ elements and writing back $O(n^{1/c}\log n)$ elements and clean-up phase not being able to reduce this to $n^{1/c}$ until the last level of the recursion. One can prevent this by observing that when distributing $n^{2/c}$ elements among $n^{1/c}$ output buckets, every bucket will have at most $n^{1/c}$ of its elements in every output bucket, with very high probability (by using Chernoff bounds [19, Chapter 4]). Hence, after distributing $O(\log n)$ elements from buckets of size $n^{1/c}$, we can make another sequential pass, reading $n^{1/c}\log n$ elements at a time (i.e., the elements that were contributed by batches of size $\log n$ from $n^{1/c}$ buckets) and applying the fact that all together they could not contribute more than $O(n^{1/c})$ elements, and hence writing

back only $O(n^{1/c})$. Note that we remain data-oblivious, since again we are using the observation of what the adversary would expect to see with very high probability.

We note that each recursive level does not depend on higher levels and, hence, has the same memory requirements as a single shuffle pass of the corresponding algorithm.

Theorem 6.1: *Given an integer constant $c \geq 3$ and an input array of size n, the optimized Melbourne shuffle executes $O(cn^{(c-1)/c})$ operations, each exchanging a message of size $O(n^{1/c})$ between a user with private memory of size $O(n^{1/c})$ and a server with storage of size $O(n)$. Also, the user and server perform $O(cn)$ work.*

7 Applications

In this section, we show how the Melbourne shuffle can be efficiently parallelized on a PRAM. We also show how to use the Melbourne shuffle to obtain an efficient oblivious storage solution.

7.1 PRAM

In the EREW PRAM model, $\sqrt{n}$ processors can agree on a seed for a permutation of the shuffle and run the shuffle as specified below. If privacy is not an issue, encryption and decryption calls can be omitted from the algorithms to get a data independent shuffle algorithm.

Theorem 7.1: *The (optimized) Melbourne shuffle can be executed in $O(1)$ steps by $\sqrt{n}$ processors, each with $O(\sqrt{n}\log n)$ ($O(\sqrt{n})$) private memory, that access a shared memory of size $O(n \log n)$ ($O(n)$) via EREW protocol.*

Proof: (Sketch) Label each processor with an id from 0 to $\sqrt{n} - 1$. During the shuffle pass in Algorithm 2 every bucket id_I of I can be read by the id_Ith processor from the shared memory. The processor creates batches of elements that go from his bucket to output buckets in the shared memory. He then can write the batches to the locations of the shared memory of O as specified in line 20 of Algorithm 2 by using processor's id instead of id_I. During the clean-up phase each processor can read the bucket of O with the same id as his processor id and write it back. (Changes to Algorithm 4 are analogous.) ■

7.2 Oblivious Storage

In this section, we give an overview of a secure and efficient oblivious storage scheme that uses the Melbourne shuffle. The oblivious storage (OS) we consider here follows the framework proposed in [6] and the follow-up work of [12]. The goal of the oblivious storage is to hide client's access pattern to his remotely stored data from anyone observing it, including the storage provider. Informally, OS transforms a virtual sequence of requests into a simulated one that appears to be data-independent. This is achieved by a mixture of accesses that are the same for every access sequence (e.g., the Melbourne shuffle) and of accesses that are randomized, and come from the same distribution, hence they appear to be independent of the access sequence.

OS Scheme Our OS scheme consists of *setup*, *access* and *rebuild* phases. The setup phase arranges, encrypts and outsources the data to the remote storage server. The access phase transforms a virtual request into a sequence of accesses to the remote storage. Once these accesses are performed, the requested element is returned. After a batch of requests, the data at the server is shuffled in order to be able to proceed with the access phase for the next batch.

We provide a short description of the *square root OS solution* [6] and show how the Melbourne Shuffle improves its performance.

In the following description we assume that the user and the server can exchange messages of size $O(\sqrt{n})$ and the client's memory is also $O(\sqrt{n})$.

Setup Let A be an array on n items. The client extends A by adding $\sqrt{n}$ fake elements with keys $n + 1, n + 2, \ldots, n + \sqrt{n}$. He then encrypts A using Enc-IND-CPA secure encryption to get an array I, which he sends to the server to store. The client picks a secret permutation, PRP π and calls $\mathsf{shuffle}(I, \pi, O)$ where O is the location at the server where the shuffled and encrypted array of A is stored. We set $I \leftarrow O$ for the access phase. The user also allocates at the server an empty cache C that can fit encryptions of up to $\sqrt{n}$ requested elements. As we shall see, the number of non-empty locations in C is the total number of requests that were made to remote storage since the last time I was shuffled. We refer to how many elements are present in C as l. The client remembers the seed that generated π so that he can find the elements during the access phase.

Access Phase Given the cache C, the encrypted array I and a request $\mathsf{read}(x)$ or $\mathsf{write}(x, v')$ the access phase creates the following oblivious simulation that depends on n, the total number of accesses made so far, and secret PRP π. It starts by reading the cache with a single request since the cache fits in one request message and in client's memory. It decrypts the cache and checks if an element with the key x is present in C. If so, it remembers the element (x, v). Then, a location in I is accessed as follow. If the element was found in the cache a fake element with the key $n + l$ is accessed by requesting the location $\pi(n + l)$ of I. Otherwise, the location that stores the element with the key x is accessed, by requesting location $\pi(x)$ of I. After reading the cache and making one request to I, the user has the desired element (x, v). If the original request was read he writes encrypted element (x, v) to the first empty location in C on the server, if it was a write, the client writes (x, v') instead. This phase can proceed this way for $\sqrt{n} - 1$ more requests, after that the cache fills up and the rebuild phase follows. This phase requires 3 accesses to the remote storage per every requested element.

Rebuild Phase The goal of the rebuild phase is to free C by placing updated elements back to I and shuffle I using a new secret permutation π' that is independent of π. This step has to be done in a data-oblivious manner to prevent correlations between access patterns to I before and after reshuffle. One first writes C to I by reading C in private memory and then reading buckets of size $\sqrt{n}$ of I, updating them with elements of C, if needed, and writing them back re-encrypted. After merging the cache C and I we are ready to call the Melbourne shuffle via $\mathsf{shuffle}(I, \pi', O)$. The client updates the seed that generated π' in his private memory and allocates an empty cache C at the server. The rebuild takes $O(\sqrt{n})$ accesses if we use the optimized Melbourne Shuffle.

Deamortized Oblivious Storage The overall cost of the shuffle can be amortized over a batch of $\sqrt{n}$ requests to achieve $O(1)$ overhead. We could also deamortize this by using the method of [11]. The method involves doubling the space at the server by adding a second cache of size $\sqrt{n}$ and memory of size $n + \sqrt{n}$ where the rebuild is happening "behind the scenes" during the access phase. The rebuild on the new array is done by batching its requests with the requests from the access phase. The client now does a constant increase in the amount of work and has a new permutation ready when the cache is full. The performance of the above OS scheme base on the optimized Melbourne shuffle is summarized in the following theorem.

Theorem 7.2: *The randomized oblivious storage scheme based on the optimized Melbourne shuffle has the following properties, where n is the size of the outsourced dataset:*

- *The private memory at the client and each message exchanged between the client and server have size $O(\sqrt{n})$.*
- *The memory at the server has size $O(n)$.*
- *The access overhead to perform a storage request is $O(1)$.*

Extension to Small Messages The recursion could be applied to square root solution, when messages of size $n^{1/c}$ are used to exchange between the user and the server, and the user has $n^{1/c}$ private memory. The solution contains one cache of size $n^{1/c}$, that is small enough to fit into private memory and be read using one request, and $c-1$ levels. Each level i is large enough to contain $n^{(i+1)/c}$ real elements and $n^{i/c}$ fake elements. The cache and levels 1 to $i-1$ have a similar cache functionality for level i as the cache C in the square root solution, except together they can store $O(n^{i/c})$ previously accessed elements. Each level $i < c-1$ contains $n^{(i+1)/c}$ buckets of size $O(\log n)$ that allows one to store $n^{(i+1)/c}$ elements in a hash table and avoid collisions with very high probability. Buckets that have less than $\log n$ elements, real or fake,are filled in with dummies. The last level, level $c-1$, has n elements, hence a permutation can be used to store and access the elements.

Given the memory arrangement of the extension with small messages at the server from Section 7.2, the access and rebuild phases proceed as follows.

Access phase requires a total of $c+1$ accesses to the server: read the cache of size $n^{1/c}$, read $O(\log n)$ size bucket from every $c-2$ levels, read one element from the last level, write an updated cache back. Note that each bucket can be read into memory since we assume private memory and messages of size $O(n^{1/c})$. Hence, $O(c)$ accesses are required to access an element obliviously. The access phase proceeds as follows: read the cache, if an element is found access a new fake element from level 1, and proceed with fake accesses from then on. Otherwise, look up the bucket using a hash function of level 1 where the element is supposed to be if it was read before. Again, proceed with consequent fake accesses, if the element was found, or keep looking for the element.

Rebuild phase: After $n^{1/c}$ elements have been accessed, level 1 is rebuilt taking $O(n^{1/c}\log n)$ accesses to be rebuilt using the Melbourne Shuffle for small messages (see Section 6). Similarly, when (i-1)th level is full it requires shuffling of $n^{(i+1)/c}\log n$ elements at level i using $O(i \times n^{i/c}\log n)$ accesses. Finally, the last level requires $O((c-1)n^{(c-1)/c})$ accesses to rebuild. We could amortize the cost of the rebuilds to get $O(c\log n)$ amortized, or deamortized using [11], access overhead per every element. We could also increase message size to $O(\sqrt[c]{n}\log n)$ and use the optimized Melbourne shuffle to achieve a constant overhead from the rebuild, since now the rebuild for level i takes $O(n^{i/c})$ accesses and can be (de)amortized over $O(n^{i/c})$ elements that caused it. Hence, with messages of size $\sqrt[c]{n}\log n$, and using the optimized Melbourne shuffle we get $O(c)$ (de)amortized access overhead.

Though, this solution resembles the Hierarchical Solution of [6] by using buckets of size $\log n$ at every level, it has two important differences. The expansion factor from level to level is $n^{1/c}$, instead of 2, hence, we only have c levels, and the rebuild phase uses our shuffle algorithm instead of a more expensive oblivious sort.

Theorem 7.3: *The randomized oblivious storage scheme based on the Melbourne shuffle with small messages has the following properties, where n is the size of the outsourced dataset and c is a constant such that $c \geq 3$:*

- *The private memory at the client and each message exchanged between the client and server have size $O(\sqrt[c]{n})$ ($O(\sqrt[c]{n}\log n)$).*

- *The memory at the server has size $O(n)$.*
- *The access overhead to perform a storage request is $O(c \log n)$ $(O(c))$.*

Acknowledgements

This research was supported in part by the National Science Foundation under grants CNS–1011840, CNS–1012060, CNS–1228485, CNS–1228639, and IIS–124758, by the National Institutes of Health under grant R01-CA180776, and by the Office of Naval Research under grant N00014-08-1-1015. Olga Ohrimenko worked on this project in part while at Brown University.

References

[1] M. Ajtai, J. Komlós, and E. Szemerédi. An $O(n \log n)$ sorting network. In *ACM Symp on Theory of Computing (STOC)*, pages 1–9, 1983.

[2] D. Aldous and P. Diaconis. Shuffling cards and stopping times. *The American Mathematical Monthly*, 93(5):333–348, 1986.

[3] K. E. Batcher. Sorting networks and their applications. In *Proc. 1968 Spring Joint Computer Conf.*, pages 307–314. AFIPS Press, 1968.

[4] S. Chen, R. Wang, X. Wang, and K. Zhang. Side-channel leaks in web applications: a reality today, a challenge tomorrow. In *Proc. IEEE Symposium on Security and Privacy*, pages 191–206, 2010.

[5] O. Goldreich. Towards a theory of software protection and simulation by oblivious RAMs. In *ACM Symp. on Theory of Computing (STOC)*, pages 182–194, 1987.

[6] O. Goldreich and R. Ostrovsky. Software protection and simulation on oblivious RAMs. *J. ACM*, 43(3):431–473, 1996.

[7] M. T. Goodrich. Randomized Shellsort: A simple oblivious sorting algorithm. In *Proc. ACM-SIAM Sump. on Discrete Algorithms (SODA)*, pages 1–16, 2010.

[8] M. T. Goodrich. Zig-zag Sort: A Deterministic Data-Oblivious Sorting Algorithm Running in $O(n \log n)$ Time. In *ACM Symp. on Theory of Computing (STOC)*, 2014. To appear.

[9] M. T. Goodrich and M. Mitzenmacher. Privacy-Preserving Access of Outsourced Data via Oblivious RAM Simulation. In *Proc. Int. Colloq. on Automata, Languages and Programming (ICALP)*, pages 576–587, 2011.

[10] M. T. Goodrich and M. Mitzenmacher. Anonymous card shuffling and its applications to parallel mixnets. In *ICALP (2)*, pages 549–560, 2012.

[11] M. T. Goodrich, M. Mitzenmacher, O. Ohrimenko, and R. Tamassia. Oblivious RAM simulation with efficient worst-case access overhead. In *Proc. ACM Workshop on Cloud Computing Security (CCSW)*, pages 95–100, 2011.

[12] M. T. Goodrich, M. Mitzenmacher, O. Ohrimenko, and R. Tamassia. Practical oblivious storage. In *ACM Conf. on Data and Application Security and Privacy (CODASPY)*, pages 13–24, 2012.

[13] M. T. Goodrich, M. Mitzenmacher, O. Ohrimenko, and R. Tamassia. Privacy-preserving group data access via stateless oblivious RAM simulation. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 157–167, 2012.

[14] M. S. Islam, M. Kuzu, and M. Kantarcioglu. Access pattern disclosure on searchable encryption: Ramification, attack and mitigation. In *NDSS*, 2012.

[15] S. Kamara and K. Lauter. Cryptographic cloud storage. In *Financial Cryptography and Data Security*, volume 6054 of *LNCS*, pages 136–149. Springer, 2010.

[16] J. Katz and Y. Lindell. *Introduction to Modern Cryptography*. Chapman and Hall/CRC Press, 2007.

[17] D. E. Knuth. *The Art of Computer Programming, Vol. 2: Seminumerical Algorithms*. Addison-Wesley Longman, 3 edition, 1997.

[18] E. Kushilevitz, S. Lu, and R. Ostrovsky. On the (in)security of hash-based oblivious RAM and a new balancing scheme. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 143–156, 2012.

[19] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.

[20] B. Pinkas and T. Reinman. Oblivious RAM revisited. In T. Rabin, editor, *Advances in Cryptology (CRYPTO)*, volume 6223 of *LNCS*, pages 502–519. Springer, 2010.

[21] E. Shi, T.-H. H. Chan, E. Stefanov, and M. Li. Oblivious RAM with $O((\log N)^3)$ worst-case cost. In *Proc. ASIACRYPT*, pages 197–214, 2011.

[22] E. Stefanov, E. Shi, and D. Song. Towards Practical Oblivious RAM. In *Proc. Network and Distributed System Security Symposium (NDSS)*, 2012.

[23] E. Stefanov, M. van Dijk, E. Shi, C. Fletcher, L. Ren, X. Yu, and S. Devadas. Path ORAM: An Extremely Simple Oblivious RAM Protocol. In *ACM Conf. on Computer and Communications Security*, CCS, 2013.

[24] E. O. Thorp. Nonrandom shuffling with applications to the game of Faro. *Journal of the American Statistical Association*, 68(344):pp. 842–847, 1973.

[25] P. Williams and R. Sion. Single round access privacy on outsourced storage. In *ACM Conf. on Computer and Communications Security (CCS)*, pages 293–304, 2012.

[26] P. Williams, R. Sion, and B. Carbunar. Building castles out of mud: practical access pattern privacy and correctness on untrusted storage. In *ACM Conference on Computer and Communications Security (CCS)*, pages 139–148, 2008.